

# Parsimonious and Efficient Likelihood Composition by Gibbs Sampling

Davide Ferrari and Guoqi Qian

Department of Mathematics and Statistics, University of Melbourne

## Abstract

The traditional maximum likelihood estimator (MLE) is often of limited use in complex high-dimensional data due to the intractability of the underlying likelihood function. Maximum composite likelihood estimation (McLE) avoids full likelihood specification by combining a number of partial likelihood objects depending on small data subsets, thus enabling inference for complex data. A fundamental difficulty in making the McLE approach practicable is the selection from numerous candidate likelihood objects for constructing the composite likelihood function. In this paper, we propose a flexible Gibbs sampling scheme for optimal selection of sub-likelihood components. The sampled composite likelihood functions are shown to converge to the one maximally informative on the unknown parameters in equilibrium, since sub-likelihood objects are chosen with probability depending on the variance of the corresponding McLE. A penalized version of our method generates sparse likelihoods with a relatively small number of components when the data complexity is intense. Our algorithms are illustrated through numerical examples on simulated data as well as real genotype SNP data from a case-control study.

Keywords: Composite likelihood estimation, Gibbs sampling, Jackknife, Efficiency, Model Selection

## 1. INTRODUCTION

While maximum likelihood estimation plays a central role in statistical inference, today its application is challenged in a number of fields where modern technologies allow scientists to collect data in unprecedented size and complexity. These fields include genetics, biology, environmental research, meteorology and physics, to name a few. Two main issues arise when attempting to apply traditional maximum likelihood to high-dimensional or complex data. The first concerns modelling and model selection, since high-dimensional data typically imply complex models for which the full likelihood function is difficult, or impossible, to specify. The second relates to computing, when the full likelihood function is available but it is just too complex to be evaluated.

The above limitations of traditional maximum likelihood have motivated the development of composite likelihood methods, which avoid the issues from full likelihood maximization by combining a set of low-dimensional likelihood objects. Besag (1974) was an early proponent of composite likelihood estimation in the context of data with spatial dependence; Lindsay (1988) developed composite likelihood inference in its generality and systematically studied its properties. Over the years, composite likelihood methods have proved useful in a range of complex applications, including models for geostatistics, spatial extremes and statistical genetics. See Varin, Reid and Firth (2011) for a comprehensive survey of composite likelihood theory and applications.

Let  $\mathbf{X} = (X_1, \dots, X_d)^T$  be a  $d \times 1$  random vector with pdf (or pmf)  $f(\mathbf{x}|\boldsymbol{\theta}_0)$ , where  $\boldsymbol{\theta}_0 \in \boldsymbol{\Theta} \subseteq \mathbb{R}^p$ ,  $p \geq 1$ , is the unknown parameter. From independent observations  $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)}$ , one typically computes the maximum likelihood estimator (MLE),  $\hat{\boldsymbol{\theta}}_{mle}$ , by maximizing the likelihood function  $L_{mle}(\boldsymbol{\theta}) = \prod_{i=1}^n f(\mathbf{X}^{(i)}|\boldsymbol{\theta})$ . Now suppose that complications in the  $d$ -dimensional pdf (pmf)  $f(\mathbf{x}|\boldsymbol{\theta})$  make it difficult to specify (or compute)  $L_{mle}(\boldsymbol{\theta})$  as the data dimension  $d$  grows, but it is relatively easy to specify (or compute) one-, two-, ..., dimensional distributions up to some order for some functions of  $X_1, \dots, X_d$ . One can then follow Lindsay (1988) to estimate  $\boldsymbol{\theta}$  by the maximum composite likelihood estimator (McLE), which

maximizes the composite likelihood function:

$$L_{cl}(\boldsymbol{\theta}) = \prod_{m=1}^M L_m(\boldsymbol{\theta}), \quad (1)$$

where each  $L_m(\boldsymbol{\theta})$  is a user-specified partial likelihood (or sub-likelihood) depending on marginal or conditional events on variables. For example,  $L_m$  can be defined using a marginal event  $\{x_m\}$  (marginal composite likelihood), pairs of variables such as  $\{x_1, x_2\}$  (pairwise likelihood), or conditional events like  $\{x_1, x_2\}|\{x_1\}$  (conditional composite likelihood). For simplicity, we assume that  $\boldsymbol{\theta}$  is common to all sub-likelihood components, so that any factorization based on a subset of  $\{L_m(\boldsymbol{\theta}), m = 1, \dots, M\}$  yields a valid objective function.

Although the composite likelihood approach provides a flexible framework with a sound theory for making inference about  $\boldsymbol{\theta}$  in situations involving multivariate data, there exist at least two challenges hindering the efficiency improvement and feasible computing of McLE in applications. The first challenge lies with selecting the right sub-likelihood components for constructing an informative composite likelihood function. The current practice of keeping all plausible factors in (1) is not well justified in terms of efficiency relative to MLE, since inclusion of redundant factors can deteriorate dramatically the variance of the corresponding composite likelihood estimator (e.g., see Cox and Reid (2004)). A better strategy would be to choose a subset of likelihood components which are maximally informative on  $\boldsymbol{\theta}_0$ , and drop noisy or redundant components to the maximum extent. However, little work is found in the literature in regard to optimal selection of sub-likelihood components. The second challenge lies with the computational complexity involved in the maximization of  $L_{cl}(\boldsymbol{\theta})$ , which can go quickly out of reach as  $d$  (and  $M$ ) increases. Particularly, note that computing  $L_{cl}$  involves  $M \times N_{ops}(d_{cl})$  operations, where  $N_{ops}(d_{cl})$  is the number of operations for each sub-likelihood component. The computational burden is exacerbated when  $\boldsymbol{\Theta}$  is relatively large and the different sub-likelihood factors  $L_m(\boldsymbol{\theta})$  do not depend on distinct elements of  $\boldsymbol{\theta}$ . One would like to see this computing complexity be alleviated to a manageable level by applying parsimonious likelihood composition in the presence of high (or ultra-high)

dimensional data.

Motivated by the aforementioned challenges, in this paper we propose a new class of stochastic selection algorithms for optimal likelihood composition. Our method uses Gibbs sampler, a specific Markov Chain Monte Carlo (MCMC) sampling scheme to generate informative – yet parsimonious – solutions. The resulting estimates will converge to the one maximally informative on the target parameter  $\theta_0$  as the underlying Markov chain reaches equilibrium. This is because sub-likelihood components generated in the algorithms are drawn according to probabilities determined by a criterion minimizing the McLE’s variance or its consistent approximation. Theory of unbiased estimating equations prescribes McLE’s asymptotic variance as an optimality criterion, i.e., the  $O_F$ -optimality criterion, see Heyde (1997, Ch. 2), but such an ideal objective has scarce practical utility due to the mathematical complexity and computational cost involved in evaluating the numerous covariance terms implied by the asymptotic variance expression (Lindsay, Yi and Sun 2011). To address this issue, we replace the asymptotic variance by a rather inexpensive jackknife variance estimator computed by a one-step Newton-Raphson iteration. Such a replacement is shown to work very well based on our numerical study.

Another advantage of our approach is that proper use of the Gibbs sampler can generate sparse composition rules, i.e., composite likelihoods involving a relatively small number of informative sub-likelihood components. Note that the model space implied by the set of all available sub-likelihood components can be large, even when  $d$  is moderate. For example, if  $d = 20$ , we have  $2^M = 2^{\binom{d}{2}} = 2^{190}$  possible composition rules based on pair-wise likelihood components alone. To cope with such a high-dimensionality, we combine Gibbs sampling with a composite likelihood stochastic stability mechanism. Analogously to the stochastic stability selection proposed by Meinshausen and Bühlmann (2010) in the context of high-dimensional model selection, our approach selects a small but sufficient number of informative likelihood components through the control of the error rates of false discoveries.

The paper is organized as follows: In Section 2, we outline the main framework and basic concepts related to composite likelihood estimation. We also describe the  $O_F$ -optimality

criterion and introduce its jackknife approximation. In Section 3, we describe our core algorithm for simultaneous likelihood estimation and selection. In Section 4, we discuss an extension of our algorithm by incorporating the ideas of model complexity penalty and stochastic stability selection. This leads to a second algorithm for parsimonious likelihoods composition for high-dimensional data. In Section 5, we illustrate our methods through numerical examples involving simulated data and real genetic single nucleotide polymorphism (SNP) data from a breast cancer case-control study. Section 6 concludes the paper with final remarks.

## 2. SPARSE COMPOSITE LIKELIHOOD FUNCTIONS

### 2.1 Binary Likelihood Composition

Let  $\{\mathcal{A}_1, \dots, \mathcal{A}_M\}$  be a set of marginal or conditional sample sub-spaces associated with probability density functions (pdfs)  $f_m(\mathbf{x} \in \mathcal{A}_m | \boldsymbol{\theta})$ . See Varin et al. (2011) for interpretation and examples of  $\mathcal{A}_m$ . Given independent  $d$ -dimensional observations  $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)} \sim f(\mathbf{x} | \boldsymbol{\theta}_0)$ ,  $\boldsymbol{\theta}_0 \in \boldsymbol{\Theta} \subseteq \mathbb{R}^p$ ,  $p \geq 1$ , we define the composite log-likelihood function:

$$\ell_{cl}(\boldsymbol{\theta}, \boldsymbol{\omega}) = \sum_{m=1}^M \omega_m \ell_m(\boldsymbol{\theta}), \quad (2)$$

where  $\boldsymbol{\omega} = (\omega_1, \dots, \omega_M)^T \in \boldsymbol{\Omega} = \{0, 1\}^M$ , and  $\ell_m(\boldsymbol{\theta})$  is the partial log-likelihood

$$\ell_m(\boldsymbol{\theta}) = \sum_{i=1}^n \log f_m(\mathbf{X}^{(i)} \in \mathcal{A}_m | \boldsymbol{\theta}). \quad (3)$$

Each partial likelihood object (sub-likelihood)  $\ell_m(\boldsymbol{\theta})$  is allowed to be selected or not, depending on whether  $\omega_m$  is 1 or 0. We aim to approximate the unknown complete log-likelihood function  $\ell_{mle}(\boldsymbol{\theta}) = \log L_{mle}(\boldsymbol{\theta})$  by selecting a few – yet the most informative – sub-likelihood objects from the  $M$  available objects, where  $M$  is allowed to be larger than  $n$  and  $p$ . Given the composition rule  $\boldsymbol{\omega} \in \boldsymbol{\Omega}$ , the maximum composite likelihood estimator (McLE), denoted

by  $\hat{\boldsymbol{\theta}}(\boldsymbol{\omega})$ , is defined by the solution of the following system of estimating equations:

$$\mathbf{0} = \sum_{i=1}^n \mathbf{U}^{(i)}(\boldsymbol{\theta}, \boldsymbol{\omega}) := \sum_{i=1}^n \sum_{m=1}^M \omega_m \mathbf{U}_m^{(i)}(\boldsymbol{\theta}), \quad (4)$$

where  $\mathbf{U}_m^{(i)}(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \log f_m(\mathbf{X}^{(i)} \in \mathcal{A}_m | \boldsymbol{\theta})$ ,  $m = 1, \dots, M$ ,  $i = 1, \dots, n$ , are unbiased partial scores functions. Under standard regularity conditions on the sub-likelihoods (Lindsay 1988),  $\sqrt{n}(\hat{\boldsymbol{\theta}}(\boldsymbol{\omega}) - \boldsymbol{\theta}_0(\boldsymbol{\omega})) \xrightarrow{\mathcal{D}} N_p(\mathbf{0}, \mathbf{V}_0)$  with asymptotic variance given by the  $p \times p$  matrix

$$\mathbf{V}_0(\boldsymbol{\omega}) = \mathbf{V}(\boldsymbol{\theta}_0, \boldsymbol{\omega}) = \mathbf{H}(\boldsymbol{\theta}_0, \boldsymbol{\omega})^{-1} \mathbf{K}(\boldsymbol{\theta}_0, \boldsymbol{\omega}) \mathbf{H}(\boldsymbol{\theta}_0, \boldsymbol{\omega})^{-1}, \quad (5)$$

where

$$\mathbf{H}(\boldsymbol{\theta}, \boldsymbol{\omega}) = \sum_{m=1}^M \omega_m \mathbf{H}_m(\boldsymbol{\theta}), \quad \mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\omega}) = \text{Var} \left[ \sum_{m=1}^M \omega_m \mathbf{U}_m(\boldsymbol{\theta}) \right],$$

$\mathbf{H}_m(\boldsymbol{\theta}) = \text{Var}(\mathbf{U}_m(\boldsymbol{\theta}))$  is the  $p \times p$  sensitivity matrix for the  $m$ th component, and  $\mathbf{U}_m(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \log f_m(\mathbf{X} \in \mathcal{A}_m | \boldsymbol{\theta})$ .

The main approach followed here is to minimize, in some sense, the asymptotic variance,  $\mathbf{V}_0(\boldsymbol{\omega})$ . To this end, theory of unbiased estimating equations suggests that both matrices  $\mathbf{H}$  and  $\mathbf{K}$  should be considered in order to achieve this goal (e.g., see Heyde (1997, Chapter 2)). On one hand, note that  $\mathbf{H}$  measures the covariance between the composite likelihood score  $\mathbf{U}(\boldsymbol{\theta}, \boldsymbol{\omega}) = \sum_{m=1}^M \omega_m \mathbf{U}_m(\boldsymbol{\theta}) = \sum_{m=1}^M \omega_m \nabla_{\boldsymbol{\theta}} \log f_m(\mathbf{X} \in \mathcal{A}_m | \boldsymbol{\theta})$  and the MLE score  $\mathbf{U}_{mle}(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \log f(\mathbf{X}; \boldsymbol{\theta})$ . To see this, differentiate both sides in  $E[\mathbf{U}(\boldsymbol{\theta}_0, \boldsymbol{\omega})] = \mathbf{0}$  and obtain  $\mathbf{H}(\boldsymbol{\theta}_0, \boldsymbol{\omega}) = E[\mathbf{U}_{mle}(\boldsymbol{\theta}_0) \mathbf{U}(\boldsymbol{\theta}_0, \boldsymbol{\omega})^T]$ . This shows that adding sub-likelihood components is desirable, since it increases the covariance with the full likelihood. On the other hand, including too many correlated sub-likelihoods components inflates the variance through the covariance terms in  $\mathbf{K}(\boldsymbol{\theta}_0, \boldsymbol{\omega})$ .

## 2.2 Fixed-Sample Optimality and its Jackknife Approximation

The objective of minimizing the asymptotic variance is still undefined since  $\mathbf{V}(\boldsymbol{\theta}_0, \boldsymbol{\omega})$  in (5) is a  $p \times p$  positive semidefinite matrix. Therefore, we consider the following one-dimensional

objective function

$$g_0(\boldsymbol{\omega}) = \log \det\{\mathbf{V}(\boldsymbol{\theta}_0, \boldsymbol{\omega})\} = \log \det\{\mathbf{K}(\boldsymbol{\theta}_0, \boldsymbol{\omega})\} - 2 \log \det\{\mathbf{H}(\boldsymbol{\theta}_0, \boldsymbol{\omega})\}. \quad (6)$$

The minimizer,  $\boldsymbol{\omega}_0$ , of the ideal objective (6) corresponds to the  $O_F$ -optimal solution (fixed-sample optimality) (e.g., see Heyde (1997), Chapter 2). Clearly, such a program still lacks practical relevance, since (6) depends on the unknown parameter  $\boldsymbol{\theta}_0$ . Therefore,  $g_0(\boldsymbol{\omega})$  should be replaced by some sample-based estimate, say  $\hat{g}_0(\boldsymbol{\omega})$ .

One option is to use the following consistent estimates of  $\mathbf{H}(\boldsymbol{\theta}_0, \boldsymbol{\omega})$  and  $\mathbf{K}(\boldsymbol{\theta}_0, \boldsymbol{\omega})$  in (6):

$$\hat{\mathbf{H}}(\boldsymbol{\omega}) = \frac{1}{n-1} \sum_{m=1}^M \sum_{i=1}^n \omega_m \mathbf{U}_m^{(i)}(\hat{\boldsymbol{\theta}}) \mathbf{U}_m^{(i)}(\hat{\boldsymbol{\theta}})^T, \quad \hat{\mathbf{K}}(\boldsymbol{\omega}) = \frac{1}{n} \sum_{i=1}^n \mathbf{U}^{(i)}(\hat{\boldsymbol{\theta}}, \boldsymbol{\omega}) \mathbf{U}^{(i)}(\hat{\boldsymbol{\theta}}, \boldsymbol{\omega})^T, \quad (7)$$

where the estimator  $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}(\boldsymbol{\omega})$  is the McLE. Although this strategy works in simple models when  $n$  is relatively large and  $M$  is small, the estimator  $\hat{\mathbf{K}}(\boldsymbol{\omega})$  is knowingly unstable when  $n$  is small compared to  $\dim(\boldsymbol{\Theta})$  (Varin et al. 2011). Another issue with this approach in high-dimensional datasets is that the number of operations required to compute  $\hat{\mathbf{K}}(\boldsymbol{\omega})$  (or  $\mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\omega})$ ) increases quadratically in  $\sum_{m=1}^M \omega_m$ .

To reduce the computational burden and avoid numerical instabilities, we estimate  $g_0$  by the following one-step jackknife criterion:

$$\hat{g}(\boldsymbol{\omega}) = \log \det \left\{ \sum_{i=1}^n \left( \hat{\boldsymbol{\theta}}(\boldsymbol{\omega})^{(-i)} - \bar{\boldsymbol{\theta}}(\boldsymbol{\omega}) \right) \left( \hat{\boldsymbol{\theta}}(\boldsymbol{\omega})^{(-i)} - \bar{\boldsymbol{\theta}}(\boldsymbol{\omega}) \right)^T \right\}, \quad (8)$$

where the pseudo-value  $\hat{\boldsymbol{\theta}}(\boldsymbol{\omega})^{(-i)}$  is a composite likelihood estimator based on a sample without observation  $\mathbf{X}^{(i)}$ , and  $\bar{\boldsymbol{\theta}}(\boldsymbol{\omega}) = \sum_{i=1}^n \hat{\boldsymbol{\theta}}(\boldsymbol{\omega})^{(-i)} / n$ . Alternatively, one could use the delete- $k$  jackknife estimate where  $k > 1$  observations at the time are deleted to compute the pseudo-values. The delete- $k$  version is computationally cheaper than delete-1 jackknife and therefore should be preferred when the sample size,  $n$ , is moderate or large. Other approaches – including bootstrap – should be considered depending on the model set up.

For example, block re-sampling techniques such as the block-bootstrap (see Hall, Horowitz and Jing (1995) and subsequent papers) are viable options for spatial data and time series.

When the sub-likelihood scores are in closed form, the pseudo-values can be efficiently approximated by the following one-step Newton-Raphson iteration:

$$\hat{\boldsymbol{\theta}}(\boldsymbol{\omega})^{(-i)} = \tilde{\boldsymbol{\theta}} + \left( \sum_{m=1}^M \omega_m \sum_{j \neq i} \mathbf{U}_m^{(j)}(\tilde{\boldsymbol{\theta}}) \mathbf{U}_m^{(j)}(\tilde{\boldsymbol{\theta}})^T \right)^{-1} \left( \sum_{m=1}^M \omega_m \sum_{j \neq i} \mathbf{U}_m^{(j)}(\tilde{\boldsymbol{\theta}}) \right), \quad (9)$$

where  $\tilde{\boldsymbol{\theta}}$  is any root- $n$  consistent estimator of  $\boldsymbol{\theta}$ . Note that  $\tilde{\boldsymbol{\theta}}$  does not need to coincide with the McLE,  $\hat{\boldsymbol{\theta}}(\boldsymbol{\omega})$ , so a computationally cheap initial estimator – based only on a few small sub-likelihoods subset – may be considered. Remarkably, the number of operations required in the Newton-Raphson iteration (9) grows linearly in the number of sub-likelihood components, so that the one-step jackknife objective has computational complexity comparable to a single evaluation of the composite likelihood function (4). Large sample properties of jackknife estimator of McLE’s asymptotic variance can be derived under regularity conditions on  $\mathbf{U}_m$  and  $\mathbf{H}_m$  analogous to those described in Shao (1992). Then, for the one-step jackknife using the root- $n$  consistent starting point  $\tilde{\boldsymbol{\theta}}$ , we have  $n(\hat{g}(\boldsymbol{\omega}) - g_0(\boldsymbol{\omega})) \xrightarrow{p} 0$ , as  $n \rightarrow \infty$ , uniformly on  $\boldsymbol{\Omega}$ . Moreover, the estimator  $\hat{g}(\boldsymbol{\omega})$  is asymptotically equivalent to the classic jackknife estimator.

### 3. PARAMETER ESTIMATION AND LIKELIHOOD COMPOSITION

#### 3.1 Likelihood Composition via Gibbs Sampling

When computing the McLE of  $\boldsymbol{\theta}$ , our objective is to find an optimal binary vector  $\boldsymbol{\omega}^*$  to estimate  $\boldsymbol{\omega}_0 = \operatorname{argmin}_{\boldsymbol{\omega} \in \boldsymbol{\Omega}} g_0(\boldsymbol{\omega})$ , where  $g_0(\cdot)$  is the ideal objective function defined in (6). Typically, the population quantity  $g_0$  cannot be directly assessed, so we replace  $g_0$  with the sample-based jackknife estimate  $\hat{g}$  described in Section 2 and aim at finding

$$\boldsymbol{\omega}^* = \operatorname{argmin}_{\boldsymbol{\omega} \in \boldsymbol{\Omega}} \hat{g}(\boldsymbol{\omega}).$$



This task, however, is computationally infeasible through enumerating the space  $\Omega$  if  $d$  is even moderately large. For example, for composite likelihoods defined based on all pairs of variables  $(X_s, X_t)$ ,  $1 \leq s < t \leq d$  (pairwise likelihood),  $\Omega$  contains  $2^{\binom{d}{2}} = 2^{190}$  elements when  $d = 20$ .

To overcome this enumeration complexity, we carry out a random search method based on Gibbs sampling. We regard the weight vector  $\omega$  as a random vector following the joint probability mass function (pmf)

$$\pi_\tau(\omega) = \frac{1}{Z(\tau)} \exp\{-\tau \hat{g}(\omega)\}, \quad \omega \in \Omega, \quad (10)$$

where  $Z(\tau) = \sum_{\omega \in \Omega} \exp\{-\tau \hat{g}(\omega)\}$  is the normalizing constant. The above distribution depends on the tuning parameter  $\tau > 0$ , which controls the extent to which we emphasize larger probabilities (and reduce smaller probabilities) on  $\Omega$ . Then  $\omega^*$  is also the mode of  $\pi_\tau(\omega)$ , meaning that  $\omega^*$  will have the highest probability to appear, and will be more likely to appear earlier rather than later, if a random sequence of  $\omega$  is to be generated from  $\pi_\tau(\omega)$ .

Therefore, estimating  $\omega^*$  can be readily done based on the random sequence generated from  $\pi_\tau(\omega)$ . But generating a random sample from  $\pi_\tau(\omega)$  directly is difficult because  $\pi_\tau(\omega)$  contains an intractable normalizing constant  $Z(\tau)$ . Instead, we will generate a Markov chain using the product of all univariate conditional pmf's with respect to  $\pi_\tau(\omega)$  as the transitional kernel. The stationary distribution of such a Markov chain can be proved to be  $\pi_\tau(\omega)$  (Robert and Casella 2004, Chapter 10). Hence, the part of this Markov chain after reaching equilibrium can be regarded as a random sample from  $\pi_\tau(\omega)$  for most purposes. The MCMC method just described is in fact the so-called Gibbs sampling method. The key factor for the Gibbs sampling to work is all the univariate conditional probability distributions of the target distribution can be relatively easily simulated.

Let us write  $\omega = (\omega_1, \dots, \omega_M)$ ;  $\omega_{m_1:m_2} = (\omega_{m_1}, \omega_{m_1+1}, \dots, \omega_{m_2})$ , if  $m_1 \leq m_2$ , and  $\omega_{m_1:m_2} = \emptyset$  otherwise; and  $\omega_{-m} = (\omega_1, \dots, \omega_{m-1}, \omega_{m+1}, \dots, \omega_M)$ . Then it is easy to see that

the conditional pmf of  $\omega_m$  given  $\boldsymbol{\omega}_{-m}$ ,  $m = 1, \dots, M$ , is

$$\pi_\tau(\omega_m | \boldsymbol{\omega}_{-m}) = \frac{\exp\{-\tau \hat{g}(\boldsymbol{\omega})\}}{\exp\{-\tau \hat{g}(\boldsymbol{\omega}_m^{[0]})\} + \exp\{-\tau \hat{g}(\boldsymbol{\omega}_m^{[1]})\}}, \quad \omega_m = 0, 1, \quad (11)$$

where  $\omega_m^{[j]} = (\omega_{1:(m-1)}, j, \omega_{(m+1):M})$ ,  $j = 0, 1$ . Note that  $\pi_\tau(\omega_m | \boldsymbol{\omega}_{-m})$  is simply a Bernoulli pmf and so it is easily generated. For each binary vector  $\boldsymbol{\omega}$ ,  $\hat{g}(\boldsymbol{\omega})$  is computed using the one-step jackknife estimator described in Section 2. Therefore, the probability mass function  $\pi_\tau(\boldsymbol{\omega})$  is well defined for any  $\tau > 0$ . On the other hand, we have shown that the Gibbs sampling can be used to generate a Markov chain from  $\pi_\tau(\boldsymbol{\omega})$ , from which we can find a consistent estimator  $\hat{\boldsymbol{\omega}}^*$  for the mode  $\boldsymbol{\omega}^*$  if  $\boldsymbol{\omega}^*$  is unique. Consequently, the universally maximum composite log-likelihood estimator of  $\boldsymbol{\theta}$  can be approximated by the McLE,  $\hat{\boldsymbol{\theta}}(\hat{\boldsymbol{\omega}}^*)$ .

Note that the mode  $\boldsymbol{\omega}^*$  is not necessarily unique. In this case one can still consider consistent estimation for  $\boldsymbol{\omega}^*$  but in the following meaning. We know  $\hat{g}(\boldsymbol{\omega}^*)$  is the minimum and always unique according to its definition. Thus  $\hat{g}(\boldsymbol{\omega}^*)$  can be consistently and uniquely estimated based on the Markov chain of  $\hat{g}(\boldsymbol{\omega})$  induced from the Markov chain of  $\boldsymbol{\omega}$  generated from  $\pi_\tau(\boldsymbol{\omega})$  when the length of the Markov chain goes to infinity. Therefore, any estimator  $\hat{\boldsymbol{\omega}}^*$  such that  $\hat{g}(\hat{\boldsymbol{\omega}}^*)$  becomes consistent with  $\hat{g}(\boldsymbol{\omega}^*)$  can be regarded as a consistent estimator of  $\boldsymbol{\omega}^*$ . Consequently, the McLE  $\hat{\boldsymbol{\theta}}(\hat{\boldsymbol{\omega}}^*)$  for each such  $\hat{\boldsymbol{\omega}}^*$  is still the universally McLE but not necessarily unique. From a practitioner's view point there is no need to identify all consistent estimators of  $\boldsymbol{\omega}^*$  and all universally McLE of  $\boldsymbol{\theta}$ . Finding out one such  $\hat{\boldsymbol{\omega}}^*$  or a tight superset of it would be a sufficient advance in parsimonious and efficient likelihood composition.

### 3.2 Algorithm 1: MCMC Composite Likelihood Selection (MCMC-CLS)

The above discussion motivates the steps of our core Gibbs sampling algorithm for simultaneous composite likelihood estimation and selection. Let  $\tau$  be given and fixed.

0. For  $t = 0$ , choose an initial binary vector of weights  $\boldsymbol{\omega}^{(0)}$  and compute the one-step jackknife estimator  $\hat{g}(\boldsymbol{\omega}^{(0)})$ . E.g. randomly set 5 elements of  $\boldsymbol{\omega}^{(0)}$  to 1 and the rest to 0.

1. For each  $t = 1, \dots, T$  for a given  $T$ , obtain  $\boldsymbol{\omega}^{(t)}$  by repeating 1.1 to 1.4 for each  $m = 1, \dots, M$  sequentially.

1.1. Compute, if not available yet,  $\hat{g}(\omega_{1:(m-1)}^{(t)}, j, \omega_{(m+1):M}^{(t-1)})$ ,  $j = 0, 1$ .

1.2. Compute the conditional pmf of  $\omega_m$  given  $(\omega_{1:(m-1)}^{(t)}, \omega_{(m+1):M}^{(t-1)})$ :

$$\pi_\tau(\omega_m = j | \omega_{1:(m-1)}^{(t)}, \omega_{(m+1):M}^{(t-1)}) \propto \exp\{-\tau \hat{g}(\omega_{1:(m-1)}^{(t)}, j, \omega_{(m+1):M}^{(t-1)})\} \quad (12)$$

where  $j = 0, 1$ . Note this is a Bernoulli pmf.

1.3. Generate a random number from the Bernoulli pmf obtained in 1.2, and denote the result as  $\omega_m^{(t)}$ .

1.4. Set  $\boldsymbol{\omega}^{(t)} \leftarrow (\omega_{1:(m-1)}^{(t)}, \omega_m^{(t)}, \omega_{(m+1):M}^{(t-1)})^T$ . Also compute and record  $\hat{g}(\boldsymbol{\omega}^{(t)})$ .

2. Compute  $\hat{\boldsymbol{\omega}}^* = \arg \min_{1 \leq t \leq T} \hat{g}(\boldsymbol{\omega}^{(t)})$ , and regard it as the estimate of  $\boldsymbol{\omega}^*$ . Alternatively, column-combine  $\boldsymbol{\omega}^{(1)}, \dots, \boldsymbol{\omega}^{(T)}$  generated in Step 1 into an  $M \times T$  matrix  $\hat{\mathbf{W}}$ ; then compute row averages of  $\hat{\mathbf{W}}$ , say  $\bar{\omega}_1, \dots, \bar{\omega}_M$ , and set  $\tilde{\boldsymbol{\omega}}^* = (\tilde{\omega}_1^*, \dots, \tilde{\omega}_M^*)$  where  $\tilde{\omega}_m^* = 1$ , if  $\bar{\omega}_m \geq \xi$ , and  $\tilde{\omega}_m^* = 0$  if  $\bar{\omega}_m < \xi$ , where  $\xi$  is some constant larger than 0.5.

3. Finally, compute  $\hat{\boldsymbol{\theta}}(\hat{\boldsymbol{\omega}}^*)$  (or  $\hat{\boldsymbol{\theta}}(\tilde{\boldsymbol{\omega}}^*)$ ) and  $\hat{g}(\hat{\boldsymbol{\omega}}^*)$  (or  $\hat{g}(\tilde{\boldsymbol{\omega}}^*)$ ).

Firstly, note that Gibbs sampling has been used in various contexts in the literature of model selection. George and McCulloch (1993) used a similar strategy to generate the distribution of the variable indicators in Bayesian linear regression. Qian (1999) used Gibbs sampler for selecting robust linear regression models. Qian and Field (2002) used the Gibbs sampler for selecting generalized linear regression models. Brooks, Friel and King (2003) and Qian and Zhao (2007) used Gibbs sampling for selection in the context of time series models. However, to our knowledge this is the first work proposing a general-purpose Gibbs sampler for construction of composite likelihoods.

Secondly, the sequence  $\boldsymbol{\omega}^{(1)}, \dots, \boldsymbol{\omega}^{(T)}$  is a Markov chain, which requires an initial vector  $\boldsymbol{\omega}^{(0)}$  and a burn-in period to be in equilibrium. Values of  $\boldsymbol{\omega}^{(0)}$  do not affect the eventual

attainment of equilibrium so can be arbitrarily chosen. From a computational point of view most components of  $\boldsymbol{\omega}^{(0)}$  should be set to 0 to reduce computing load. For example, we can randomly set all but 5 of the components to 0. To assess whether the chain has reached equilibrium, we suggest the control-chart method discussed in Qian and Zhao (2007). For the random variable  $\hat{g}(\boldsymbol{\omega})$ , we have the following probability inequality for any given  $b > 1$ :

$$Pr\left(\hat{g}(\boldsymbol{\omega}) - \min \hat{g}(\boldsymbol{\omega}) \geq b\sqrt{Var[\hat{g}(\boldsymbol{\omega})] + (E[\hat{g}(\boldsymbol{\omega})] - \min \hat{g}(\boldsymbol{\omega}))^2}\right) \leq \frac{1}{b^2}$$

This inequality can be used to find an upper control limit for  $\hat{g}(\boldsymbol{\omega})$ . For example, by setting  $b = \sqrt{10}$ , an at least 90% upper control limit for  $\hat{g}(\boldsymbol{\omega})$  can be estimated as  $\hat{g}^* + \sqrt{10s^2 + 10(\bar{g} - \hat{g}^*)^2}$ , where  $\hat{g}^*$ ,  $\bar{g}$  and  $s^2$  are the minimum, sample mean and sample variance based on the first  $N$  observations,  $\hat{g}(\boldsymbol{\omega}^{(1)}), \dots, \hat{g}(\boldsymbol{\omega}^{(N)})$ ,  $N < T$ , where typically we set  $N = \lfloor T/2 \rfloor$ . We then count the number of observations passing the upper control limit in the remaining sample  $\hat{g}(\boldsymbol{\omega}^{(N+1)}), \dots, \hat{g}(\boldsymbol{\omega}^{(T)})$ . If more than 10% of the points are above the control limit, then at a significance level not more than 90% there is statistical evidence against equilibrium. Upper control limits of different levels for  $\hat{g}(\boldsymbol{\omega})$  can be similarly calculated and interpreted by choosing different values of  $b$ .

Thirdly,  $\hat{\boldsymbol{\omega}}^*$  computed in Step 2 is simply a sample mode of  $\pi_\tau(\boldsymbol{\omega})$  based on its definition in (10). Hence, by the Ergodic Theorem for stationary Markov chain,  $\hat{\boldsymbol{\omega}}^*$  is a strongly consistent estimator of  $\boldsymbol{\omega}$  under minimal regularity conditions. With similar arguments  $\bar{\omega}_m$  is a strongly consistent estimator of the success probability involved in the marginal distribution of  $\omega_m$  induced from  $\pi_\tau(\boldsymbol{\omega})$ . Hence, it is not difficult to see that the resultant estimator  $\tilde{\boldsymbol{\omega}}^*$  should satisfy  $\tilde{\omega}_m^* \geq \omega_m^*$  without requiring  $T$  to be very large. Propositions 1 and 2 in Qian and Zhao (2007) provide an exposition of this property. Therefore, the estimator  $\tilde{\boldsymbol{\omega}}^*$  captures all informative sub-likelihood components with high probability.

Finally, the tuning constant  $\tau$  adjusts the mixing behavior of the chain, which has important consequences on the exploration/exploitation trade-off on the search space  $\boldsymbol{\Omega}$ . If  $\tau$  is too small, the algorithm produces solutions approaching the global optimal value  $\boldsymbol{\omega}^*$  slowly; if  $\tau$

is large, then the algorithm finds local optima and may not reach  $\omega^*$ . The former behavior corresponds to a rapidly mixing chain, while the latter occurs when the chain is mixing too slowly. In the composite likelihood selection setting, the main hurdle is the computational cost, so  $\tau$  should be set according to the available computing capacity, after running some graphical or numerical diagnostics (e.g., see Robert and Casella (2004)). We choose to use  $\tau = d$  in our empirical study, which does not seem to create adverse effects.

## 4. AN EXTENSION FOR HIGH-DIMENSIONAL DATA

### 4.1 Sparsity-enforcing penalization

Without additional modifications, Algorithm 1 ignores the likelihood complexity, since solutions with many sub-likelihoods have in principle the same chance to occur as those with fewer components. To discourage selection of overly complex likelihoods, we augment the Gibbs distribution (10) as follows:

$$\pi_{\tau,\lambda}(\omega) = Z(\tau, \lambda)^{-1} \exp\{-\tau \hat{g}_\lambda(\omega)\}, \quad (13)$$

where

$$\hat{g}_\lambda(\omega) = \hat{g}(\omega) + \text{pen}(\omega|\lambda), \quad \tau > 0, \lambda > 0, \quad (14)$$

$\hat{g}(\omega)$  is the jackknifed variance objective defined in Section 2,  $Z(\tau, \lambda)$  is the normalization constant, and  $\text{pen}(\omega)$  is a complexity penalty enforcing sparse solutions when  $\dim(\Omega)$  is large. Maximization of  $\pi_{\tau,\lambda}(\omega)$  is interpreted as a maximum a posteriori estimation for  $\omega$ , where the probability distribution proportional to  $\exp\{-\text{pen}(\omega|\lambda)\}$  is regarded as a prior pmf over  $\Omega$ . In this paper, we use the penalty term of form  $\text{pen}(\omega|\lambda) = \lambda \sum_{m=1}^M \omega_m$ , since it corresponds to well-established model-selection criteria. For example, choices  $\lambda = 1$ ,  $\lambda = 2^{-1} \log n$  and  $\lambda = \log \log n$  correspond to the AIC, BIC and HQC criteria, respectively (e.g., see Claeskens and Hjort (2008)). Other penalties could be considered as well depending on the model structure and available prior information; however, these are not shown to be crucial based on our empirical study so will not be explored in this paper.

## 4.2 Composite Likelihood Stability Selection

To find the optimal solution  $\boldsymbol{\omega}^*$ , one could compute a sequence of optimal values  $\hat{\boldsymbol{\omega}}_{\lambda_1}^*, \dots, \hat{\boldsymbol{\omega}}_{\lambda_B}^*$  and then take  $\min_{1 \leq b \leq B} \hat{g}(\hat{\boldsymbol{\omega}}_{\lambda_b}^*)$ . There are, however, various issues in this approach: first, the globally optimal value  $\boldsymbol{\omega}^*$  might not be a member of the set  $\{\hat{\boldsymbol{\omega}}_{\lambda_b}^*\}_{b=1}^B$ , since the mode of  $\pi_{\tau, \lambda}(\boldsymbol{\omega})$  is not necessarily the composite likelihood solution which minimizes  $\hat{g}(\boldsymbol{\omega})$ . Second, even if  $\boldsymbol{\omega}^*$  is in such a set, determining  $\lambda$  is typically challenging. To address the above issues, we employ the idea of stability selection, introduced by Meinshausen and Bühlmann (2010) in the context of variable selection for linear models. Given an arbitrary value for  $\lambda$ , stochastic stability exploits the variability of random samples generated from  $\pi_{\tau, \lambda}(\boldsymbol{\omega})$  by the Gibbs procedure, say  $\boldsymbol{\omega}_{\lambda}^{(1)}, \dots, \boldsymbol{\omega}_{\lambda}^{(T)}$  and chooses all the partial likelihoods that occur in a large fraction of generated samples. For a given  $0 < \xi < 1$ , we define the set of stable likelihoods by the vector  $\hat{\boldsymbol{\omega}}^{\text{stable}}$ , with elements

$$\hat{\omega}_m^{\text{stable}} = \begin{cases} 1, & \text{if } \frac{1}{T} \sum_{t=1}^T \omega_{\lambda, m}^{(t)} \geq \xi, \\ 0, & \text{otherwise,} \end{cases} \quad (15)$$

so we regard as stable those sub-likelihoods selected more frequently and disregard sub-likelihood items with low selection probabilities. Following Meinshausen and Bühlmann (2010), we choose the tuning constant  $\xi$  using the following bound on the expected number of false selections,  $V$ :

$$E(V) \leq \frac{1}{(2\xi - 1)} \frac{\eta_{\lambda}}{M}, \quad (16)$$

where  $\eta_{\lambda}$  is average number of selected sub-likelihood components. In multiple testing, the quantity  $\alpha = E(V)/M$  is sometimes called the per-comparison error rate (PCER). By increasing  $\xi$ , only few likelihood components are selected, so that we reduce the expected number of falsely selected variables. We choose the threshold  $\xi$  by fixing the PCER at some desired value (e.g.,  $\alpha = 0.10$ ), and then choose  $\xi$  corresponding to the desired error rate. The unknown quantity  $\eta_{\lambda}$  in our setting can be estimated by the average number of sub-likelihood

components over  $T$  Gibbs samples.

Finally, note that tuning  $\xi$  according to (16) makes redundant the determination of the optimal  $\lambda$  value as long as  $\text{pen}(\boldsymbol{\omega}|\lambda)$  in (14) is not dominant over  $\hat{g}(\boldsymbol{\omega})$ . This is further supported by our empirical study where we found the effect of  $\lambda$  on  $\hat{\boldsymbol{\omega}}^{\text{stable}}$  is negligible.

#### 4.3 Algorithm 2: MCMC Composite Likelihood with Stability Selection (MCMC-CLS2)

The preceding discussions lead to Algorithm 2 which is essentially the same as Algorithm 1 with two exceptions: (i) we replace  $\hat{g}$  in Algorithm 1 by the augmented objective function  $\hat{g}_\lambda$  defined in (14); (ii) Steps 2 in Algorithm 1 is replaced by the following stability selection step.

- 2'. Column-combine  $\boldsymbol{\omega}^{(1)}, \dots, \boldsymbol{\omega}^{(T)}$  generated in Step 1 into an  $M \times T$  matrix  $\hat{\mathbf{W}}$ . Compute the row averages of  $\hat{\mathbf{W}}$ , denoted as  $(\hat{\omega}_1, \dots, \hat{\omega}_M)$ , i.e.  $\hat{\omega}_m = T^{-1} \sum_{t=1}^T \omega_m^{(t)}$ ,  $m = 1, \dots, M$ . Then set  $\hat{\boldsymbol{\omega}}^{\text{stable}} = (\hat{\omega}_1^*, \dots, \hat{\omega}_M^*)$  where  $\hat{\omega}_m^* = 1$  if  $\hat{\omega}_m \geq \hat{\xi}_\lambda$  and  $\hat{\omega}_m^* = 0$  if  $\hat{\omega}_m < \hat{\xi}_\lambda$ ,  $m = 1, \dots, M$ , where  $\hat{\xi}_\lambda = \frac{1}{2} \left( \frac{\hat{\eta}_\lambda}{\alpha M^2} + 1 \right)$ , and  $\alpha$  is a nominal level for per-comparison error rate (e.g., 0.05 or 0.1).

The estimated threshold  $\hat{\xi}_\lambda$  is obtained from (16) by plugging-in  $\hat{\eta}_\lambda = \sum_{t=1}^T \sum_{m=1}^M \omega_m^{(t)} / T$ , the sample average of the numbers of selected sub-likelihood components in  $T$  Gibbs samples.

## 5. NUMERICAL EXAMPLES

### 5.1 Normal Variables with Common Location

Let  $\mathbf{X} \sim N_d(\mu \mathbf{1}, \boldsymbol{\Sigma})$ , where the parameter of interest is the common location  $\mu$ . We study the scenario where many components bring redundant information on  $\mu$  by considering covariance matrix  $\boldsymbol{\Sigma}$  with elements  $\{\boldsymbol{\Sigma}\}_{mm} = 1$ , for all  $1 \leq m \leq d$ , and off-diagonal elements  $\{\boldsymbol{\Sigma}\}_{lm} = \rho > 0$  if  $l, m \leq d^*$ , for some  $d^* < d$ , while  $\{\boldsymbol{\Sigma}\}_{lm} = 0$  elsewhere.

We consider one-wise score composite likelihood estimator solving  $0 = \sum_{m=1}^d \omega_m U_m(\mu)$ , where  $U_m(\mu) = \sum_{i=1}^n (X_m^{(i)} - \mu)$ . It is easy to find the composite likelihood estimator is  $\hat{\mu}_{cl}(\boldsymbol{\omega}) = \sum_{m=1}^d \omega_m \bar{X}_m / \sum_{m=1}^d \omega_m$  where  $\bar{X}_m = \sum_{i=1}^n X_m^{(i)} / n$ . For this simple model, the jackknife criterion  $\hat{g}(\cdot)$  can be easily computed in closed form. The pseudo-values are

$\hat{\mu}_{cl}^{(-i)}(\boldsymbol{\omega}) = \sum_{m=1}^d \omega_m \bar{X}_m^{(-i)} / \sum_{m=1}^d \omega_m$ , and the average of pseudo-values is  $\hat{\mu}_{cl}$ . It can be shown that

$$\hat{g}(\boldsymbol{\omega}) = \log \sum_{i=1}^n \left( \sum_{m=1}^d \omega_m (X_m^{(i)} - \bar{X}_m) \right)^2 - 2 \log \left( \sum_{m=1}^d \omega_m \right),$$

up to a constant not depending on  $\boldsymbol{\omega}$ . It can also be shown that  $O_F$ -criterion has the following expression

$$g_0(\boldsymbol{\omega}) = \log \text{Var}(\hat{\mu}_{cl}(\boldsymbol{\omega})) \propto \log \left( \sum_{m=1}^d \omega_m + 2\rho \sum_{l < m \leq d^*} \omega_l \omega_m \right) - 2 \log \left( \sum_{m=1}^d \omega_m \right), \quad (17)$$

depending on the unknown parameter  $\rho$ . This suggests that including too many correlated (redundant) components (with  $\rho \neq 0$ ) damages McLE's variance as  $d$  increases. Particularly, setting  $\omega_j = 1$ , for all  $j = 1, \dots, d$ , implies  $\text{Var}(\hat{\mu}_{cl}(\boldsymbol{\omega})) = O(1)$ , while choosing only uncorrelated sub-likelihoods ( $\omega_j = 0$  if  $2 \leq j \leq d^*$ , and  $\omega_j = 1$  elsewhere), gives  $\text{Var}(\hat{\mu}_{cl}(\boldsymbol{\omega})) = O(d^{-1})$ .

In Table 1, we show Monte Carlo simulation results from  $B = 250$  runs of Algorithm 1 (MCMC-CLS1) using the two approaches described in Section 3.2: one consists of choosing the best weights vector (CLS1 min); the other uses thresholding, i.e. selects elements when the weights are selected with a sufficiently large frequency (CLS1 thres.). In the same table we also report results on the Algorithm 2 (MCMC-CLS2) based on the stochastic stability selection approach. Our algorithms are compared with the estimator including all one-wise sub-likelihood components (No selection) and the optimal maximum likelihood estimator (MLE). We compute the MLE of  $\mu$  in two ways: either based on using the known  $\Sigma$  value or based on using the sample covariance  $\hat{\Sigma} = (n-1)^{-1} \sum_{i=1}^n (X_i - \bar{X})(X_i - \bar{X})^T$ . Note that the latter is not available when  $d > n$ . Note our algorithms are designed to obtain simultaneous estimation and dimension reduction in the presence of limited information (i.e.  $d > n$  or  $d \gg n$ ) where the optimal weights are difficult to estimate from the data. Stochastic selection with  $\xi = 0.7$  and  $T = 10d$  was carried out for samples of  $n = 5, 25$  and 100 observations from a model with  $d^* = 0.8 \times d$  correlated components, with  $d = 10, 30$ . To compare the



methods we computed Monte Carlo estimates of the variance ( $Var$ ) and squared bias ( $Bias^2$ ) of  $\hat{\mu}_{cl}(\omega)$ . Table 2 further reports the average number of selected likelihood components (no. comp).

For all considered data dimensions, our selection methods outperform the all one-wise likelihood (AOW or No selection) estimator and show relatively small losses in terms of mean squared error ( $Var + Bias^2$ ) compared to the MLEs. The gains in terms of variance reduction are particularly evident for larger data dimensions (e.g., see  $d = 30$ ). At the same time, our method tends to select mostly uncorrelated sub-likelihoods, while discarding the redundant components which do not contribute useful information on  $\mu$ .

Next, we illustrate the selection procedure based on MCMC-CLS2 (Algorithm 2). We draw a random sample of  $n = 50$  observations from the model  $N_d(\mu\mathbf{1}, \Sigma(\rho))$  described above with  $d = 250$ ,  $d^* = 0.8d$ , and  $\rho = 0.9$ . This corresponds to 200 strongly redundant variables and 50 independent variables. We applied Algorithm 2 with  $\alpha = 0.1$  and  $\lambda = 1$  (corresponding to the AIC penalty). In Figure 1 (b), we show the relative frequencies of the components,  $\bar{\omega}_m = \sum_{t=1}^T \omega_m^{(t)}$ ,  $m = 1, \dots, 250$ , in  $T = 1000$  MCMC iterations. As expected, the uncorrelated sub-likelihood components (components 201–250) are sampled much more frequently than the redundant ones (components 1–200). Figure 1 (c) shows the objective function  $\hat{g}_\lambda(\hat{\omega}^{\text{stable}})$  (up to a constant) evaluated at the best solution computed from past MCMC samples ( $\xi = 0.7$ ). Figure 1 (d) shows the Hamming distance,  $dist(\omega, \omega') = \sum_j I(\omega_j \neq \omega'_j)$ , between the current selected rule,  $\hat{\omega}^*$  and true optimal value,  $\omega^* = (1, \underbrace{0, \dots, 0}_{199}, \underbrace{1, \dots, 1}_{50})$ . Overall, our algorithm quickly approaches the optimal solution and the final selection has 94.0% asymptotic relative efficiency (ARE) compared to MLE. When no stochastic selection is applied and all 250 sub-likelihood components are included, the relative efficiency is only 13%. As far as computing time is concerned, our non-optimized R implementations of the MCMC-CLS1 and MCMC-CLS2 algorithms for the above example takes, respectively 4.25 and 1.87 seconds per MCMC iteration, respectively. The computing time was measured on a laptop computer with Intel ®Core™ i7-2620M CPU 2.70GHz.

$n$	$d$	$\rho$	No selection			CLS1 (min)			CLS1 (thresh.)			CLS2			MLE (known $\Sigma$ )			MLE (unknown $\Sigma$ )		
			$Var$	$Bias^2$		$Var$	$Bias^2$		$Var$	$Bias^2$		$Var$	$Bias^2$		$Var$	$Bias^2$		$Var$	$Bias^2$	
5	10	0.50	80.26	7.19		79.18	0.04		81.36	0.00		87.55	0.01		52.52	0.04		NA	NA	
		0.90	115.06	10.31		96.97	0.05		114.82	0.08		114.90	0.02		62.28	0.04		NA	NA	
	30	0.50	80.08	7.18		60.61	1.60		54.06	1.18		52.16	0.74		26.22	0.03		NA	NA	
		0.90	103.50	9.28		61.27	0.01		44.14	0.01		44.68	0.00		24.93	0.02		NA	NA	
25	10	0.50	17.14	8.28		15.27	0.18		17.45	0.06		17.73	0.06		12.01	0.03		20.99	0.03	
		0.90	19.89	1.78		14.75	0.07		18.65	0.03		19.08	0.02		13.02	0.02		21.64	0.04	
	30	0.50	11.87	1.06		7.38	0.12		6.14	0.00		6.07	0.00		4.86	0.00		NA	NA	
		0.90	24.51	2.20		11.59	0.02		6.69	0.00		6.75	0.00		5.73	0.01		NA	NA	
100	10	0.50	4.14	0.37		3.09	0.02		3.55	0.03		4.24	0.09		2.56	0.02		2.79	0.02	
		0.90	6.18	0.55		3.25	0.00		4.58	0.00		4.58	0.00		3.12	0.00		3.37	0.00	
	30	0.50	3.40	0.30		1.79	0.00		1.60	0.00		1.60	0.00		1.20	0.00		1.69	0.00	
		0.90	5.54	0.50		2.60	0.01		1.68	0.01		1.68	0.01		1.41	0.01		2.03	0.01	

Table 1: Bias and variance of estimators for the location model  $\mathbf{X} \sim N_d(\mu\mathbf{1}, \Sigma(\rho))$  by different methods. MCMC-CLS1 algorithm with and without thresholding (CLS1 (min) and MCMC-CLS1 (thresh.), respectively); MCMC-CLS2 algorithm with stability selection (CLS2) with  $\alpha = 0.1$  and  $\lambda = 1$ ; Maximum likelihood estimator based on  $\Sigma^{-1}$  where  $\Sigma$  is known; Maximum likelihood estimator based on  $\hat{\Sigma}^{-1}$ , where  $\hat{\Sigma}$  can be estimated if  $d < n$ , otherwise the value is missing and denoted by NA. For each method we show the finite sample variance ( $Var$ ) and squared bias ( $Bias^2$ ) for  $n = 5, 25, 100$ ,  $d = 10, 30$  and  $\rho = 0.5, 0.9$ . Estimates based on  $B = 250$  Monte Carlo runs (simulation settings:  $\tau = d$ ,  $T = 10d$ ,  $\xi = 0.7$ ). Monte Carlo standard errors are smaller than 0.001.

$n$	$d$	$\rho$	No selection	CLS1 (min)	CLS1 (thres.)	CLS2
5	10	0.50	10	3.77	3.92	3.72
		0.90	10	3.28	2.58	2.36
	30	0.50	30	13.94	11.01	10.84
		0.90	30	12.04	6.53	6.43
25	10	0.50	10	4.30	3.58	3.26
		0.90	10	3.27	2.04	2.00
	30	0.50	30	12.81	7.70	7.48
		0.90	30	11.96	5.98	5.98
100	10	0.50	10	4.28	2.56	2.31
		0.90	10	3.17	2.00	2.00
	30	0.50	30	12.22	6.08	6.05
		0.90	30	11.94	6.00	6.00

Table 2: Number of selected sub-likelihoods by different methods. MCMC-CLS1 algorithm with and without thresholding (CLS1 (min) and MCMC-CLS1 (thres.), respectively); MCMC-CLS2 algorithm with stability selection (CLS2) with  $\alpha = 0.1$  and  $\lambda = 1$ . For each method we show results for  $n = 5, 25, 100$ ,  $d = 10, 30$  and  $\rho = 0.5, 0.9$ . Estimates based on  $B = 250$  Monte Carlo runs (simulation settings:  $\tau = d$ ,  $T = 10d$ ,  $\xi = 0.7$ ). Monte Carlo standard errors are smaller than 0.01.

## 5.2 Exchangeable Normal Variables with Unknown Correlation

Let  $\mathbf{X} \sim N_d(\mathbf{0}, \Sigma(\rho))$ , where  $\Sigma(\rho) = \{(1 - \rho)\mathbf{I}_d + \rho\mathbf{1}_d\mathbf{1}_d^T\}$  and  $0 \leq \rho \leq 1$  is the unknown parameter of interest. The marginal univariate sub-likelihoods do not contain information on  $\rho$ , so we consider pairwise sub-likelihoods

$$\ell_{lm}(\rho) = -\frac{n}{2} \log(1 - \rho^2) - \frac{(SS_{ll} + SS_{mm})}{2(1 - \rho^2)} + \frac{\rho SS_{lm}}{1 - \rho^2}, \quad 1 \leq l < m \leq d, \quad (18)$$

where  $SS_{mm} = \sum_{i=1}^n (X_m^{(i)})^2$  and  $SS_{lm} = \sum_{i=1}^n X_l^{(i)} X_m^{(i)}$ . Given  $\boldsymbol{\omega}$ , we estimate  $\rho$  by solving the composite score equation  $0 = \sum_{j < k} \omega_{jk} U_{jk}(\rho)$  by Newton-Raphson iterations where each

pairwise score

$$U_{jk}(\rho) = (1 + \rho^2)SS_{jk} - \rho(SS_{jj} + SS_{kk}) + n\rho(1 - \rho^2) \quad (19)$$

is a cubic function of  $\rho$ . It is well known that the composite likelihood estimation can lead to poor results for this model. Cox and Reid (2004) carries out asymptotic variance calculations, showing that efficiency losses compared to MLE occur whenever  $d > 2$ , with more pronounced efficiency losses for  $\rho$  near 0.5. Particularly, if  $d = 2$  (exactly one pair),  $ARE = 1$ ; if  $d > 2$ ,  $ARE=1$  if  $\rho$  approaches 0 or 1. The next simulation results show that in finite samples composite likelihood selection is advantageous even in such a challenging situation.

Since closed-form pairwise score expressions are available for this model, we use the objective function  $\hat{g}(\boldsymbol{\omega})$  based on the one-step jackknife with pseudo-values computed as in Equation (9). In Table 3, we present results from  $B = 250$  Monte Carlo runs of the MCMC-CLS algorithm applied to the model with  $d = 5, 8, 10$  dimensions, which correspond to  $M = \binom{d}{2} = 10, 28, 45$  sub-likelihoods, respectively. We compute Monte Carlo estimates for the finite-sample relative efficiency  $RE = \widehat{Var}(\hat{\rho}_{apw}) / \widehat{Var}(\hat{\rho}_{cl}(\hat{\boldsymbol{\omega}}^*))$ , where  $\widehat{Var}$  denotes the Monte Carlo estimate of the finite-sample variance  $\hat{\rho}_{cl}(\hat{\boldsymbol{\omega}}^*)$  is the estimator selected by the MCMC-CLS algorithm, while  $\hat{\rho}_{apw}$  is the all pairwise (APW) estimator obtained by including all available pairs (thus  $RE > 1$  indicates that our stochastic selection outperforms no selection). We show values of  $\rho$  around 0.5, since they correspond to the largest asymptotic efficiency losses of pairwise likelihood compared to MLE (see Cox and Reid (2004), Figure 1). In all considered cases, our stochastic selection method improves the efficiency of the estimator based on all pairwise components; at the same time, our composite likelihoods employ a much smaller number of components. For example, when  $\rho = 0.6$  the efficiency improvements range from 8% to 39%, using only about half of the available components.

### 5.3 Real Data Analysis: Australian Breast Cancer Family Study

In this section, we apply the MCMC-CLS algorithm to a real genetic dataset of women with breast cancer obtained from the Australian Breast Cancer Family Study (ABCFS) (Dite,

		$n = 10$			$n = 50$		
$M = \binom{d}{2} =$		10	28	45	10	28	45
$\rho = 0.4$	$RE$	1.27(0.03)	1.11(0.01)	1.07(0.01)	1.15(0.01)	1.12(0.01)	1.06(0.01)
	No. comp.	5.83(0.09)	13.70(0.16)	21.12(0.22)	7.30(0.08)	13.82(0.16)	21.45(0.22)
$\rho = 0.5$	$RE$	1.36(0.02)	1.14(0.01)	1.06(0.01)	1.19(0.01)	1.11(0.01)	1.06(0.01)
	No. comp.	5.51(0.08)	13.32(0.07)	20.96(0.22)	7.01(0.07)	13.64(0.18)	21.23(0.22)
$\rho = 0.6$	$RE$	1.39(0.02)	1.17(0.01)	1.10(0.01)	1.38(0.02)	1.14(0.01)	1.08(0.01)
	No. comp.	5.80(0.08)	12.60(0.16)	20.99(0.20)	5.39(0.08)	13.27(0.15)	21.45(0.21)

Table 3: Pairwise likelihood selection for Model 2,  $N_d(\mathbf{0}, \mathbf{\Sigma}(\rho))$  based on Algorithm 1: Monte Carlo estimates for: (i) the relative efficiency of the parameter estimate under selection versus no selection (  $RE = Var(\hat{\rho}_{apw})/Var(\hat{\rho}_{cl}(\hat{\omega}^*))$ ), so that  $RE > 1$  indicates that the selection outperforms no selection); (ii) and number of sub-likelihood components (No. comp.). Monte Carlo standard errors are in parenthesis. Simulation settings:  $\tau = d$ , chain length  $T = 10d$ ,  $\xi = 0.7$ .

Jenkins, Southey, Hocking, Giles, McCredie, Venter and Hopper 2003) and control subjects from the Australian Mammographic Density Twins and Sisters Study (Odefrey, Stone, Gurrin, Byrnes, Apicella, Dite, Cawson, Giles, Treloar, English et al. 2010). All women were genotyped using a Human610-Quad beadchip array. The final data set that we used consisted of a subset of 20 single nucleotide polymorphisms (SNPs) corresponding to genes encoding a candidate susceptibility pathway, which is motivated by biological considerations. After recommended data cleaning and quality control procedures (e.g., checks for SNP missingness, duplicate relatedness, population outliers (Weale 2010)), the final data consisted of  $n = 333$  observations (67 cases and 266 controls).

To detect group effects due to cancer, we consider an extension of the latent multivariate Gaussian model first introduced by Han and Pan (2012). Let  $\mathbf{Y}^{(i)} = (Y_1^{(i)}, \dots, Y_d^{(i)})$ ,  $i = 1, \dots, n$ , be independent observations of a multivariate categorical variable measured on  $n$  subjects. Each variable  $Y_k^{(i)}$  can take values 0, 1 or 2, representing the copy number of one of

the alleles of SNP  $k$  of subject  $i$ . The binary variable  $X^{(i)} = x^{(i)} = 0$  or  $1$  represents disease status of the  $i$ th subject ( $0 = \text{control}$  and  $1 = \text{disease}$ ). We assume a latent random  $d$ -vector  $\mathbf{Z}^{(i)} = (Z_1^{(i)}, \dots, Z_d^{(i)}) \sim N_d(\boldsymbol{\mu}^{(i)}(\theta), \mathbf{R})$ , where  $\boldsymbol{\mu}^{(i)}(\theta)$  is a conditional mean vector with elements  $\mu_1^{(i)}(\theta) = \dots = \mu_d^{(i)}(\theta) = \theta x^{(i)}$  and  $\mathbf{R}$  is the correlation matrix. Our main interest is on the unknown mean parameter  $\theta$ , which is common to all the SNP variables and represents the main effect due to disease. We assume  $P(Y_k^{(i)} = 0 | X^{(i)} = x^{(i)}) = P(Z_k^{(i)} \leq \gamma_{k1})$ ,  $P(Y_k^{(i)} = 1 | X^{(i)} = x^{(i)}) = P(\gamma_{k1} < Z_k^{(i)} \leq \gamma_{k2})$ , and  $P(Y_k^{(i)} = 2 | X^{(i)} = x^{(i)}) = P(Z_k^{(i)} > \gamma_{k2})$ , where  $\gamma_{k1}$  and  $\gamma_{k2}$  are SNP-specific thresholds. The above model reflects the ordinal nature of genotypes and assumes absence of the Hardy-Weinberg equilibrium (HWE) (allele frequencies and genotypes in a population are constant from generation to generation). If the HWE holds the parameters  $\gamma_{1k}$  and  $\gamma_{2k}$  are not needed, since we have the additional constraint  $P(Y_k^{(i)} = 2) = P(Y_k^{(i)} = 1)^2$ .

Let  $\boldsymbol{\gamma} = \{(\gamma_{1k}, \gamma_{2k}) : k = 1, \dots, d\}$  and define intervals  $\Gamma_k(Y_k^{(i)})$  to be  $(-\infty, \gamma_{k1}]$ ,  $(\gamma_{k1}, \gamma_{k2}]$  and  $(\gamma_{k2}, \infty)$ , corresponding to  $Y_k^{(i)} = 0, 1$  and  $2$ , respectively. The full log-likelihood function is

$$\begin{aligned} \ell(\theta, \boldsymbol{\gamma}, \mathbf{R}) &= \sum_{i=1}^n \log P(\mathbf{Y}^{(i)} = \mathbf{y}^{(i)} | X^{(i)} = x^{(i)}) \\ &= \sum_{i=1}^n \log \int_{\Gamma_1(y_1^{(i)})} \dots \int_{\Gamma_d(y_d^{(i)})} f(z_1, \dots, z_d | \boldsymbol{\mu}^{(i)}(\theta), \mathbf{R}) dz_1 \dots dz_d, \end{aligned}$$

where  $f(z_1, \dots, z_d | \boldsymbol{\mu}, \mathbf{R})$  is the pdf of the  $d$ -variate normal density function with mean  $\boldsymbol{\mu}$  and correlation matrix  $\mathbf{R}$ . Clearly, the full log-likelihood function is intractable when  $d$  is moderate or large, due to the multivariate integral in the likelihood expression. Note that for the marginal latent components, we have  $Z_k^{(i)} \sim N_1(\theta x^{(i)}, 1)$ , so  $\boldsymbol{\gamma}$  and  $\theta$  can be estimated

by minimizing the one-wise composite log-likelihood

$$\ell_{cl}(\theta, \boldsymbol{\gamma}, \boldsymbol{\omega}) = \sum_{k=1}^d \omega_k \sum_{i=1}^n \log P(Y_k^{(i)} = y_k^{(i)} | X^{(i)} = x^{(i)}) \quad (20)$$

$$= \sum_{k=1}^d \omega_k \sum_{i=1}^n \log \int_{\Gamma_k(y_k^{(i)})} \phi(z_k | \theta x^{(i)}, 1) dz_k, \quad (21)$$

where  $\phi(\cdot | \mu, 1)$  denotes the normal pdf with mean  $\mu$  and unit variance. We focus on using the one-wise composite log-likelihood in this section, except when  $\mathbf{R}$  is to be estimated where we use pairwise composite log-likelihood. Differently from the expression in Han and Pan (2012), the disease group effect  $\theta$  is common to multiple sub-likelihood components; also, we allow for the inclusion/exclusion of particular sub-likelihood components (corresponding to SNPs) by selecting  $\boldsymbol{\omega}$ . We estimated the optimal composition rule  $\hat{\boldsymbol{\omega}}^*$  based on Gibbs samples from Algorithm 1, where the objective function  $\hat{g}(\boldsymbol{\omega}) = \log \text{Var}(\hat{\theta}(\boldsymbol{\omega}))$  was estimated by delete-10 jackknife. We selected five marginal likelihoods (SNPs) occurring with at least  $\xi = 0.7$  frequency in 250 runs of the Gibbs sampler (see Figure 2 (a)). In Figure 2 (b), we show the trajectory of the objective function  $\hat{g}(\boldsymbol{\omega})$  evaluated at the current optimal solution  $\hat{\boldsymbol{\omega}}^*$  (optimal solutions are computed from past samples using a  $\xi = 0.7$  threshold). The estimated variance tends to sway toward the minimum as more composition rules are drawn by our Gibbs sampler. This behavior is in agreement with preliminary simulation results carried out on this model (not presented here) as well as the examples presented in Sections 5.1 and 5.2.

Figure 2 (c) shows the empirical distribution of parameter estimates,  $\hat{\theta}^{(t)} = \hat{\theta}(\boldsymbol{\omega}^{(t)})$ , based on sampled vectors  $\boldsymbol{\omega}^{(t)}$ ,  $t = 1, \dots, 250$ . The vertical dashed line corresponds to the selected parameter estimate  $\hat{\theta}(\hat{\boldsymbol{\omega}}^*)$ , which is located near the mode of the empirical distribution. Particularly, the selected McLE is  $\hat{\theta}(\hat{\boldsymbol{\omega}}^*) = -0.125$  and the corresponding delete-10 jackknife standard error is  $\hat{sd}(\hat{\theta}(\hat{\boldsymbol{\omega}}^*)) = 0.012$ . The McLE based on using all 20 target SNPs is  $\hat{\theta}_{all} = -0.112$  with the delete-10 jackknife standard error  $\hat{sd}(\hat{\theta}_{all}) = 0.042$ . Our estimator gives a substantial accuracy improvement, supporting the conclusion of a

difference between case and control groups (i.e.,  $\theta \neq 0$ ) with higher confidence. Finally, in Figure 2 (d), we show estimates for the correlation matrix  $\mathbf{R}$  for the target SNPs, based on the pairwise composite likelihood described in Han and Pan (2012).

## 6. FINAL REMARKS

Composite likelihood estimation is a rapidly-growing need for a number of fields, due to the astonishing growth of data complexity and the limitations of traditional maximum likelihood estimation in this context. The Gibbs sampling protocol proposed in this paper addresses an important unresolved issue by providing a tool to automatically select the most useful sub-likelihoods from a pool of feasible components. Our numerical results on simulated and real data show that the composition rules generated by our MCMC approach are useful to improve the variance of traditional McLE estimators, typically obtained by using all sub-likelihood components available. Another advantage deriving from our method is the possibility to generate sparse composition rules, since our Gibbs sampler selects only a (relatively small) subset of informative sub-likelihoods while discarding non-informative or redundant components.

In the present paper, likelihood sparsity derives naturally from the discrete nature of our MCMC approach based on binary composition rules with  $\boldsymbol{\omega} \in \{0, 1\}^M$ . On the other hand, the development of sparsity-enforcing likelihood selection methods suited to real-valued weights would be valuable as well and could result in more efficient composition rules. For example, Lindsay et al. (2011) discuss the optimality of composite likelihood estimating functions involving both positive and negative real-valued weights. Actually, the row averages  $\bar{\omega}_1, \dots, \bar{\omega}_M$  computed in Step 2 of Algorithm 1 can also be used to replace  $\boldsymbol{\omega} = (\omega_1, \dots, \omega_M)^T$  in  $\ell_{cl}(\boldsymbol{\theta}, \boldsymbol{\omega})$  defined by (2), providing a composite log-likelihood with between-0-and-1 weights. It would be of interest to see how well this form of composite log-likelihood gets on in regard to the efficiency. This, however, was not pursued in this paper and is left for future exploration. Finally, the penalized version of the objective function described in Section 4 enforces sparse likelihood functions, which is necessary in situations



where the model complexity is relatively large compared to the sample size. Thus developing a thorough theoretical understanding on the effect of the penalty on the selection as  $d, n \rightarrow \infty$  would be very valuable for improved selection algorithm in high-dimensions.

## REFERENCES

- Besag, J. (1974), “Spatial interaction and the statistical analysis of lattice systems,” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 192–236.
- Brooks, S., Friel, N., and King, R. (2003), “Classical model selection via simulated annealing,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2), 503–520.
- Claeskens, G., and Hjort, N. L. (2008), *Model selection and model averaging*, Vol. 330 Cambridge University Press Cambridge.
- Cox, D. R., and Reid, N. (2004), “A note on pseudolikelihood constructed from marginal densities,” *Biometrika*, 91(3), 729–737.
- Dite, G. S., Jenkins, M. A., Southey, M. C., Hocking, J. S., Giles, G. G., McCredie, M. R., Venter, D. J., and Hopper, J. L. (2003), “Familial risks, early-onset breast cancer, and BRCA1 and BRCA2 germline mutations,” *Journal of the National Cancer Institute*, 95(6), 448–457.
- George, E. I., and McCulloch, R. E. (1993), “Variable selection via Gibbs sampling,” *Journal of the American Statistical Association*, 88(423), 881–889.
- Hall, P., Horowitz, J. L., and Jing, B.-Y. (1995), “On blocking rules for the bootstrap with dependent data,” *Biometrika*, 82(3), 561–574.
- Han, F., and Pan, W. (2012), “A Composite Likelihood Approach to Latent Multivariate Gaussian Modeling of SNP Data with Application to Genetic Association Testing,” *Biometrics*, 68(1), 307–315.

- Heyde, C. C. (1997), *Quasi-likelihood and its application: a general approach to optimal parameter estimation*. Springer.
- Lindsay, B. G. (1988), “Composite likelihood methods,” *Contemporary Mathematics*, 80, 221–239.
- Lindsay, B. G., Yi, G. Y., and Sun, J. (2011), “Issues and Strategies in the Selection of Composite Likelihoods,” *Statistica Sinica*, 21(1), 71–105.
- Meinshausen, N., and Bühlmann, P. (2010), “Stability selection,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4), 417–473.
- Odefrey, F., Stone, J., Gurrin, L. C., Byrnes, G. B., Apicella, C., Dite, G. S., Cawson, J. N., Giles, G. G., Treloar, S. A., English, D. R. et al. (2010), “Common genetic variants associated with breast cancer and mammographic density measures that predict disease,” *Cancer research*, 70(4), 1449–1458.
- Qian, G. (1999), “Computations and analysis in robust regression model selection using stochastic complexity,” *Computational Statistics*, 14, 293–314.
- Qian, G., and Field, C. (2002), “Using MCMC for logistic regression model selection involving large number of candidate models,” in *Monte Carlo and Quasi-Monte Carlo Methods 2000* Springer, pp. 460–474.
- Qian, G., and Zhao, X. (2007), “On time series model selection involving many candidate ARMA models,” *Computational Statistics & Data Analysis*, 51(12), 6180–6196.
- Robert, C. P., and Casella, G. (2004), *Monte Carlo statistical methods*, Vol. 319 Citeseer.
- Shao, J. (1992), “One-step jackknife for M-estimators computed using Newton’s method,” *Annals of the Institute of Statistical Mathematics*, 44(4), 687–701.
- Varin, C., Reid, N., and Firth, D. (2011), “An Overview of Composite Likelihood Methods,” *Statistica Sinica*, 21, 5–42.

Weale, M. E. (2010), “Quality control for genome-wide association studies,” in *Genetic Variation* Springer, pp. 341–372.

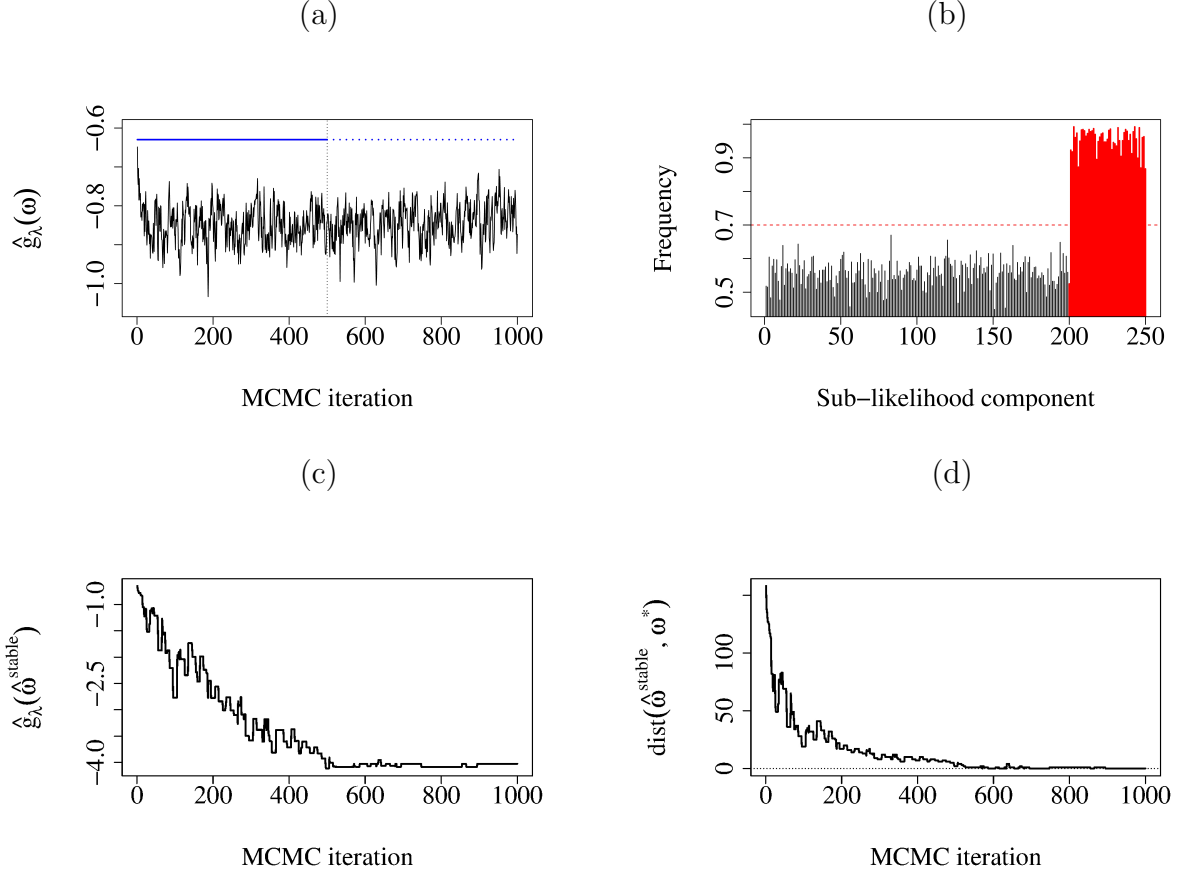


Figure 1: Stochastic selection for Model 1,  $N_d(\mu\mathbf{1}, \Sigma)$ , based on Algorithm 2. (a) Objective function  $\hat{g}_\lambda(\omega)$  evaluated at samples drawn from  $\pi_{\tau,\lambda}(\omega)$ ; the horizontal solid line is the control chart limit as described in Section 3.2. (b) Observed frequency of sampled sub-likelihood components. (c) Estimated objective function evaluated at the progressively selected likelihood,  $\hat{\omega}^{\text{stable}}$ , based on past samples. (d) Hamming distance between the progressively selected likelihood,  $\hat{\omega}^{\text{stable}}$ , and the globally optimal solution,  $\omega^*$ . Simulation settings:  $n = 50$ ,  $d = 250$ ,  $d^* = 200$ ,  $\tau = d$ ,  $T = 1000$ , burn-in length = 250,  $\alpha = 0.1$ ,  $\lambda = 1$ .

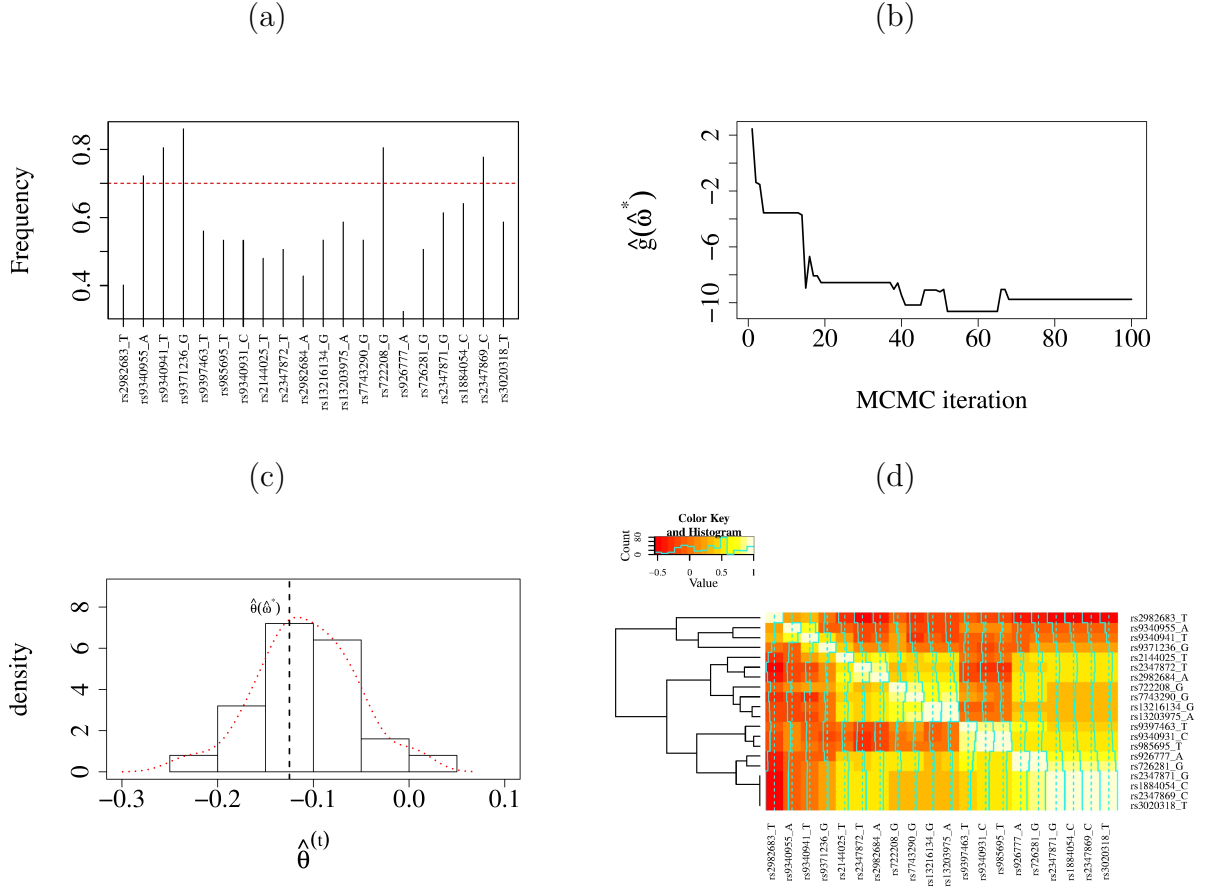


Figure 2: Composite likelihood selection for the ABCFS data by Algorithm 1: (a) Frequency of the sampled marginal likelihood components; (b) Objective function evaluated at the current solution  $\hat{\omega}^*$  (computed from past samples with  $\xi = 0.7$ ); (c) Parameter estimates,  $\hat{\theta}^{(t)} = \hat{\theta}(\omega^{(t)})$ , based on sampled composition rules,  $\omega^{(t)}$  with optimal parameter estimate  $\hat{\theta}(\omega^*)$  corresponding to vertical dashed line. (d) Pairwise likelihood estimates for the correlation matrix  $\mathbf{R}$  for SNPs in the susceptibility pathway.